AD-A247 555

‖‖‖‖‖‖‖‖‖‖‖‖‖‖‖

STARNET: an Integrated Services Broadband Optical Network with Physical Star Topology

P. Poggiolini, L. Kazovsky

Department of Electrical Engineering, Stanford University*
Durand Building, room 202, Stanford, CA 94305-4055
FAX (415) 7239251, E-mail pierluigi@star.stanford.edu

1991

DTIC
ELECTE
MAR 18 1992
D

Abstract

We propose a new broadband local area network, STAR-NET, based on a physical passive star topology. STAR-NET implements both a packet network and a high-speed WDM circuit interconnect, which are simultaneously available to every node of the network and operate independently. As a result, STARNET supports very diverse types of traffic in an optimal way. Each node requires only two lasers and its structure permits to effectively achieve frequency stabilization for the whole network. An effort toward an experimental demonstration of a 4-node, 3 Gbit/s per node, FDDI-compatible (at the packet network level) STARNET is currently in progress at Optical Communication Research Laboratory, Stanford University.

1 Introduction

Future LAN's will have to provide the wide variety of services shown in Fig. 1, [1]. The low speed services can be handled by evolutionary versions of the presently available networks. The high speed ones require a new generation of networks.

A desirable feature of a new optical broadband LAN is that it supports the whole spectrum of traffic, ranging from low-speed to high-speed and from bursty to continuous. The target is therefore an *integrated services Broadband Optical Local Area Network* (BOLAN).

Packet switching and circuit switching have both been considered for operating a BOLAN. Currently, the prevailing approach is that of packet switching, because packet switching seems to be more effective in handling

very different data rates and permits more efficient processing of bursty data. Also, virtual circuit allocation can be used to better accomodate services requiring a continuous high speed flow of information.

As far as physical implementations of BOLAN's are concerned, the passive star solution seems to be one of the most promising options, because with the use of WDM it provides a theoretically huge bandwidth. To exploit this potential in a packet switching environment, different media access strategies and node configurations have been devised. Some of these proposals rely on the use of fast tunable receivers [8, 9, 10]; others impose demanding requirements on the node ability to maintain synchronization with the other nodes and/or to operate in a time-slotted environment [3, 11, 12, 13]; still others are based on the use of multi-hop logical topologies [14, 15, 16], where each node performs some routing functions. A few of the above require that each node have many transmitters and/or receivers [13, 11].

The complexity of these proposals shows that the star topology, although very attractive from the bandwidth point of view, poses awkward problems once a packet BOLAN is to be realized on it. Essentially, the optical WDM star provides a large number of independent *circuits*. However, today's optical technology does not permit very fast switching between these optical circuits, so that exploiting them to realize a packet mode of transmission encounters hard obstacles.

In this paper we describe a BOLAN solution based on a star topology and WDM, which aims at fully exploiting the potential of this configuration but requires less hardware and protocol complexity than the above-mentioned proposals.

The paper is organized as follows.

In Section 2 we discuss the characteristics of the different types of traffic, and envisage suitable network so-

---

92 2 28 061

92-05238

‖‖‖‖‖‖‖‖‖‖‖‖‖

broadband optical local area network

Figure 1: General picture of some of today's and tomorrow's network services.

lutions. The goals of STARNET are established, and in particular the idea of implementing both a packet network and a broadband circuit interconnect on the same physical star. The rest of the paper is devoted to the presentation of the details of the proposal. In Section 3 the basic functional node structure is described. The basic network configuration is introduced. Section 4 provides an accurate description of the possible implementation alternatives for the basic node structure described in Section 3. The problem of frequency stabilization is addressed too. In Section 5, the power budget of some implementation alternatives is evaluated. In Section 6 improved versions of the packet network supported by STARNET are presented. In Section 7 an enhanced version of STARNET which also includes broadband multihop packet transport is discussed. In Section 8 broadcasting and video-conferencing using the STARNET circuit interconnect are dealt with. Section 9 describes the STARNET experiment presently under way.

2 STARNET goals

The fundamental problem in a BOLAN is that of handling different types of traffic, with very diverse speed and continuity characteristics.

We subdivide traffic into three categories:

1. low speed, bursty or continuous, including telephony;

2. high speed, continuous;

3. high speed, bursty.

Category 1 includes all the services shown in Fig. 1 as low-speed. Categories 2 and 3 share all the rest, without a sharp distinction between them: for instance, video services can be offered through transmission of a continuous stream of data, but video compression and packetization has been proposed too.

A common bus or a ring packet network is a good solution for low-speed data. If the network has sufficient capacity, telephony can be offered too on a virtual circuit basis.

For high speed communication, a circuit-switched network seems to be suitable to handle continuous streams of data, especially when the duration of the connection is relatively long. This is true for video-conferencing or transfer of big files (images, books, movies, magazines). In these instances, switching time does not need to be extremely small.

Supercomputer interconnection can be also handled on a circuit-switched basis, except when many com ers need to communicate simultaneously. In this context broadband packet communication proves necessary. This last scenario is bound to a very specific application involving a limited nuber of sources/sinks of information. Only a small subset of nodes in the LAN can potentially take full advantage of Gbit/s packet switching.

Taking into account the above observations, our pro-

posal, STARNET, offers to all nodes *both* a packet network for low-speed services *and* and a broadband circuit interconnect. These two entities operate independently and simultaneously. This arrangement will be called Fundamental Network Configuration (FNC).

As an extension to FNC, STARNET also provides a way to easily upgrade a *subset* of its nodes so that a broadband packet network can link them together without interfering with the lower speed packet network or the circuit interconnect.

As a result STARNET meets the goal of dealing with each kind of traffic in the most appropriate fashion. It also avoids the extreme constraints that other solutions impose on some network parameters or hardware features.

3 The basic node configurations

Each node of the network is optically connected to any other node through a passive optical star, as shown in Fig. 1.

Each node is equipped with one transmitter. Its light source is tunable, but tunability is required for network maintenance, flexibility and fault tolerance, rather than fast wavelength switching. In normal operation the light source is considered fixed at a certain wavelength. Each node keeps its source tuned to a different wavelength, so that a *comb of carriers* is formed.

Every node has a tunable receiver that can be dinamically tuned to any transmitter.

Using an appropriate multiplexing strategy, the node transmitter can trasmit *two independent data streams*, stream 'P' and stream 'C' as depicted in Fig. 2. This can be achieved via either TDM, or multilevel transmission or simultaneous modulation using non-interfering formats. The topic will be addressed in depth in Section 5. Letters 'P' and 'C' are mnemonic for 'Packet' and 'Circuit', and this distinction will become clear shortly.

Stream 'C' is the one which can be decoded by the tunable receivers of other nodes. Hence, apart from the multiplexing of the 'P' stream, so far we have obtained a star interconnect in which tunable receivers can tune to fixed transmitters to receive data.

In addition, each node is equipped with an *fixed receiver* (as opposed to tunable) which is capable of decoding data stream 'P' *of the previous node in the frequency comb only*.

This node structure will be referred to as Basic Node Configuration 1 (BNC1).

BNC1 nodes can be interconnected as follows. Regarding the 'P' data stream, the node can receive data from the previous node and can transmit data to the following one. Thus we can build a 'chain' along the
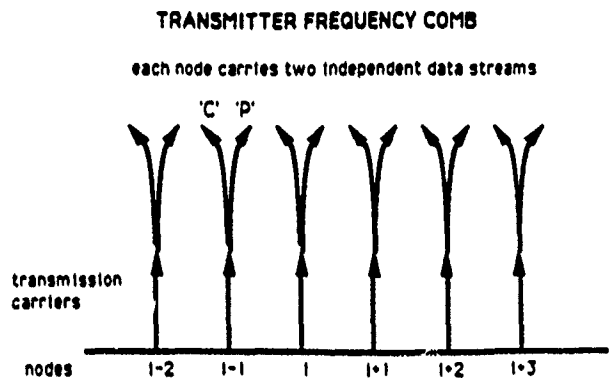

Figure 2: The node transmitter must be capable of multiplexing two independent data streams onto the same lightwave carrier.
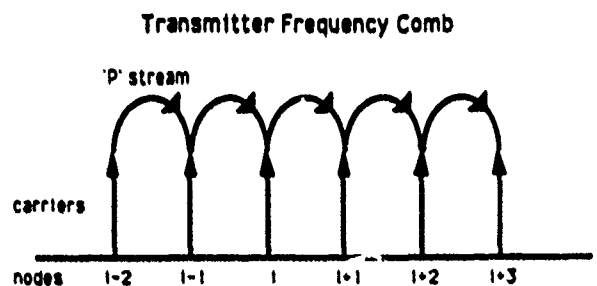

Figure 3: Each node can see the 'P' stream of the previous node. The connectivity thus achieved is depicted in figure.

frequency comb, through which information can flow as along a 'bus', provided that what is received by the previous node is relayed towards the next[1]. This situation is depicted in Fig. 3.

However, this unidirectional bus is of little use if the data stream cannot be bounced back to the top of the comb. This is accomplished as follows: the first node of the chain (the one whose transmitter is the first in the frequency comb) is equipped with a receiver that can decode the 'P' stream of the last node in the chain. Then, it relays this information down the chain, by retransmitting it through its 'P' data stream (see Fig. 4).

Topologically, we exploit this connection capability to implement a logical unidirectional ring along the frequency comb, as depicted in Fig. 5. This topology is identical to the one used by the well-known high speed packet network FDDI.

At the same time, each node has the capability to connect to any other node through the 'C' stream and its tunable receiver, which implements a circuit interconnect

---

[1] The use of the term 'bus' is somewhat improper. What we obtain here is a store-and-forward chain, but we will use 'bus' for simplicity.

Transmitter Frequency Comb

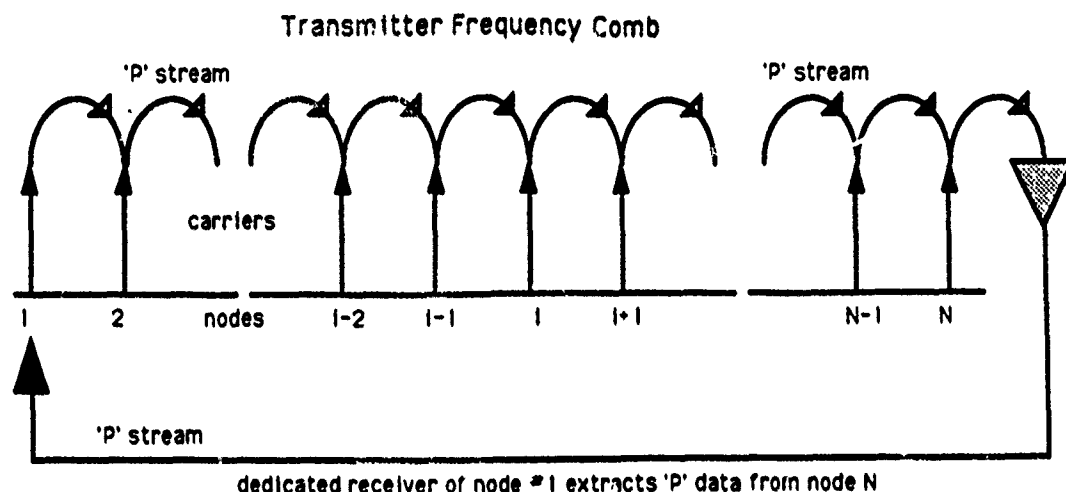dedicated receiver of node #1 extracts 'P' data from node N

Figure 4: The first node of the chain is equipped with a receiver capable of extracting the 'P' stream of the last node in the comb, and relays 'P' data downstream.
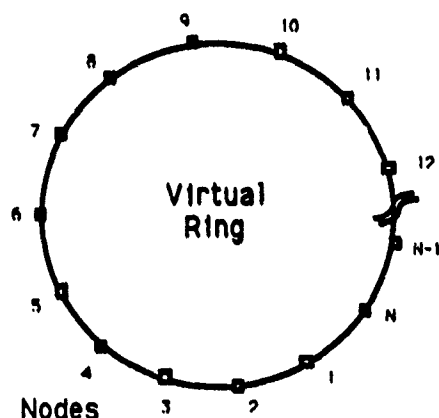

Figure 5: A logical unidirectional ring topology can be arranged on the physical underlying star network.

Figure 6: The connectivity of BNC2 permits to obtain a bidirectional bus.

among all the nodes.

The fact that the transmitters are capable of sending two multiplexed data streams not only avoids the need for the extra hardware necessary for another optical transmitter, but also prevents the formation of a second comb of carriers which would waste optical bandwidth.

An alternative to BNC1 is to equip each node with a transmitter capable of multiplexing *three* data streams, stream 'C' and streams 'P1' and 'P2'. As for BNC1, a tunable receiver can tune to any of the transmitters, and decode stream 'C'. The fixed receiver is capable to decode stream 'P1' of the previous node, in the frequency comb, and stream 'P2' of the *next* node in the frequency comb. We shall see in Section 4 that we can still use just *one* fixed receiver to accomplish this task.

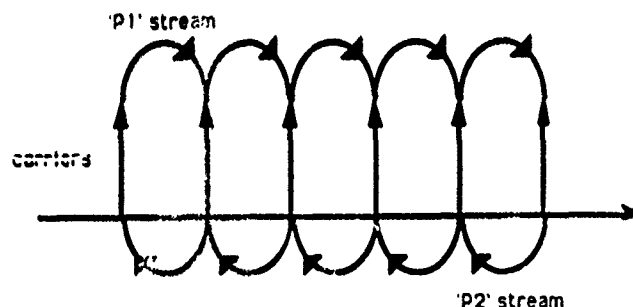This second arrangement will be called Basic Node Configuration 2 (BNC2). It permits to implement a bidi-

rectional bus (Fig. 6). If this bidirectional bus is terminated like in the previous case (Fig. 4), but at both ends, we get two independent counter-rotating rings. Otherwise, the last node in the downstream chain (see Fig. 6) can turn the data flow into the upstream chain. Similarly, the first node can relay downward the data that it receives from the upward chain. Thus, we obtain again a ring topology, but this time the stations appear twice in the ring (Fig. 7) This arrangement is attractive because there is no need for the end nodes to use special receivers to close the loop, like in Fig. 4.

Later, in Section 6, we will discuss some substantial enhancements to these topologies.

4 Node implementation

In this section we examine possible implementations for the basic node configurations.

The solutions discussed in this section are based on coherent transmission. The basic node structures are
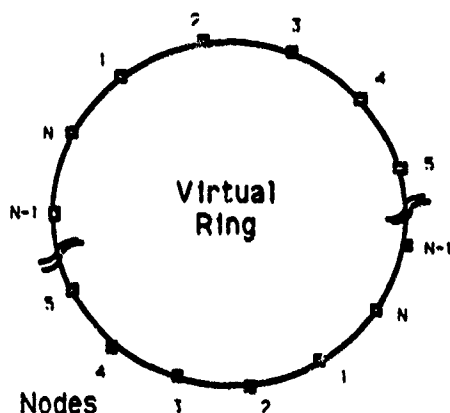
Figure 7: With BNC2 a ring topology can be arranged, where each station appears twice.

realizable with incoherent technology as well, but this topic will not be covered in this paper.

First, we consider BNC1 and propose a specific implementation in detail. Then we discuss different implementation alternatives, and extend the treatment to BNC2.

BNC1, as outlined in the previous section, has three different subsystems: a transmitter, a tunable receiver and a 'P' stream fixed receiver capable of receiving just the previous node in the frequency comb. We analyze these subsystems one by one. Fig. 8 reports a schematic of the overall node structure, to serve as a reference throughout this section.

4.1 The transmitter

The transmitter has to multiplex two independent data streams, the 'C' stream and the 'P' stream. We assume that multiplexing is done through TDM. The transmitter devotes part of the time to 'C' data and part of the time to 'P' data.

The transmission format is rather unimportant in this context. However, we assume heterodyne detection.

The transmission laser is kept at a fixed frequency in normal operation.

4.2 The tunable 'C' stream receiver

The tunable 'C' receiver is a standard heterodyne receiver with a tunable LO. In the 'C' stream TDM frame only two 'slots' are present, the 'C' one and the 'P' one. Synchronization over the desired one, 'C' in this case, can be achieved by putting a unique identifier[2] at the start of the slots. The length of both the 'P' and 'C' slots is fixed and the slots interleave periodically. As a result

synchronization is not a critical task. Since the identifier can be as short as a few bits, and the frame length can be several thousands of bits, the overhead is negligible.

The receiver laser must be tunable over the entire frequency comb. Since we are dealing with the 'C' links, which are used to support circuit-switched broadband traffic, acceptable tuning speeds could be on the order of milliseconds or even tens of milliseconds. This makes it easier to use tunable semiconductor lasers, in spite of thermal transients and mode hops.

The relaxed constraints imposed by this network arrangement could make it possible to use direct thermal tuning of DFB sources. Other types of lasers (NdYAG for instance) could be fast enough, too.

Regarding polarization problems, any means of achieving insensitivity is compatible with our network. Even polarization control could be adequate to fulfill channel switching time requirements. However, polarization control has the main drawback that in some cases an almost complete fading of the target signal may occur, so that the tuning process itself becomes critical.

Polarization diversity is an attractive solution to this problem. Rugged and low-loss polarization diversity front-ends are now appearing[3].

In addition, some new transmission systems (based for instance on polarization modulation, POLSK [17]) are insensitive to polarization fluctuations.

Alternatively, one could use polarization preserving optics in all the network.

4.3 The fixed receiver

The fixed receiver must be able to tune to the transmitter of the node that occupies the previous position in the frequency comb. This task could be carried out by another independent tunable receiver, which would remain permanently locked on the previous node carrier.

However, there is a way to build the additional receiver that does not require another laser and at the same time provides frequency stabilization of the comb of carriers.

We distinguish between two cases: in the first case an external modulator is used for transmission; in the second case, the node transmitter laser is modulated directly.

In the former case, we propose to tap part of the carrier light before it enters the modulator, and use it as LO signal for the fixed receiver. If the frequency comb of channels is arranged so that channels are equispaced, what we get at IF in the fixed receiver is depicted in Fig 9. Both the previous and the following channel in the comb are visible, but they overlap, due to the hetero-

---

[2]For instance, FDDI makes use of a very short unique identifier to mark the beginning of a packet.

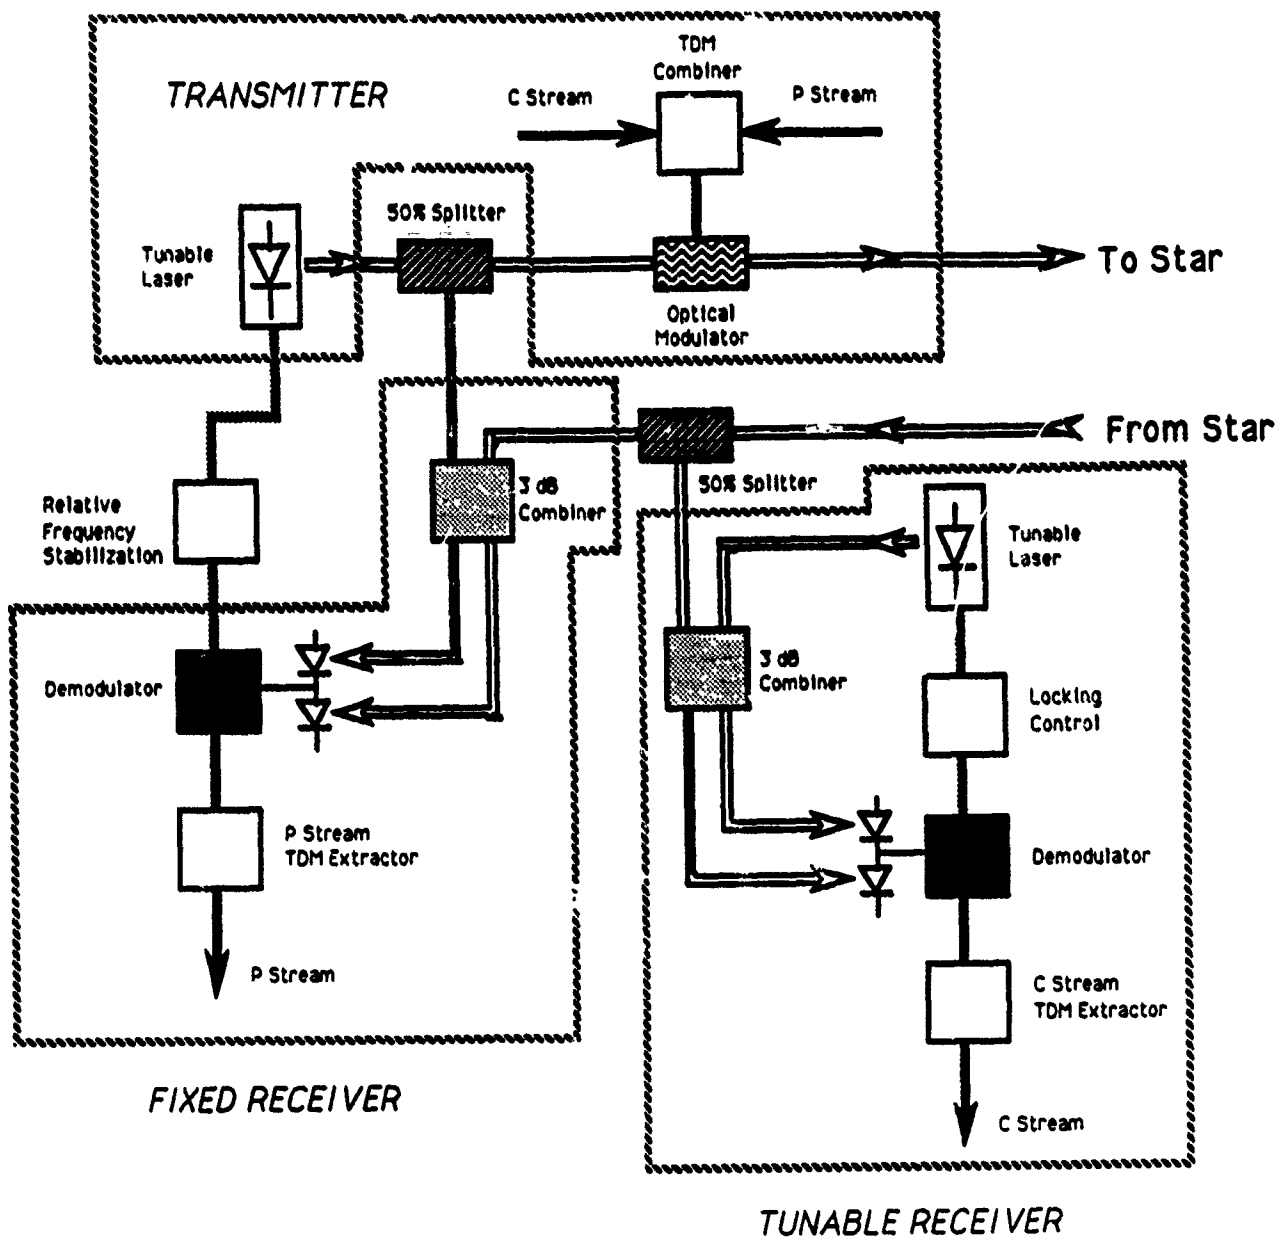[3]An all-fiber device comprising all the necessary optical processing is currently commercially available.

Figure 8: Schematic of the generic node, assuming TDM multiplexing and optical external modulation.
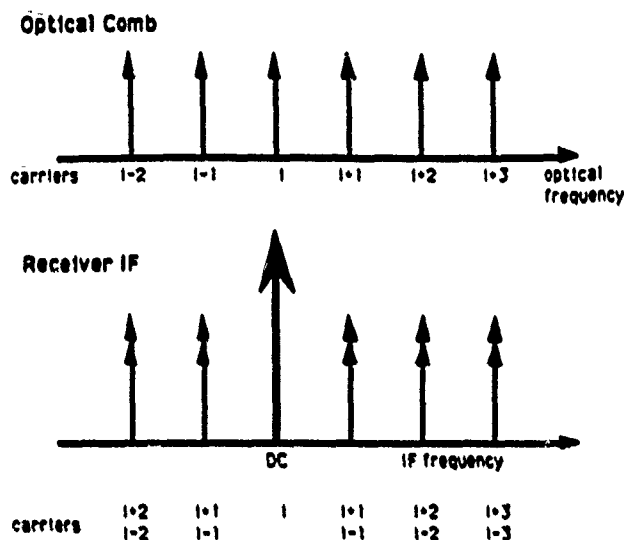
Figure 9: Carrier visibility at IF in the fixed receiver of node i, employing the node carrier as LO.

dyning process. Recalling that our objective is to decode the previous node signal, two solutions are possible.

The first one exploits image-rejection receivers. These receivers are capable of isolating either one of the two channels, suppressing the other. An interesting proposal for a balanced image-rejection receiver has been made in [4]. However the hardware overhead needed to implement it is still prohibitive.

To pave the way for a simpler solution, we resort to a new carrier allocation strategy, recently proposed in [2]. A pictorial description of it is shown in Fig. 10b, whereas in Fig. 10a the conventional equispaced frequency allocation is reported. This new concept was originally conceived to save bandwidth, and in fact it turns out that the theoretical bandwidth wasting due to heterodyne detection, with respect to homodyne or image-rejection detection, can be decreased by at least 50% as compared to the equispaced channel arrangement. The details are in [?].

Besides the bandwidth saving, another effect occurs. In Fig. 11 we show the new signal spectrum at the IF stage of the fixed receiver. The adjacent channels are now separated. and visibility over the previous node and the next node in the comb is always achievable, although the IF at which they are located maybe either one of two values. Both the previous and the next node are of interest in the second elementary node structure proposed in Section 3 (BNC2). In the simpler node structure (BNC1), one of the two signals is simply discarded.

In any case, the receiver control has to find out which is which between the two visible adjacent channels, i.e. the previous and the next one in the comb. However

this kind of identification is needed only at the network set-up time.

Through IF filters, the desired channel is isolated and sent to the demodulator, where extraction of the 'P' TDM frame is made. 'P' stream TDM frame synchronization is even easier than for the 'C' stream, because it can be done once for all at the network set-up time.

4.4 Frequency stabilization

Like any coherent receiver, the fixed receiver must keep its LO tuned so that reception of the desired channel occurs. Since the LO of the fixed receiver is also the node transmitter laser, the frequency control operated by the fixed receiver to keep its LO locked on the previous node carrier also provides relative frequency stabilization for the node carrier. Since this is done by all the nodes in the comb, each node is frequency locked to the previous node and overall network frequency stabilization is achieved.

Such idea needs experimental confirmation, since it is not clear what the effects of propagation of random frequency fluctuations along the comb could be. However, a recent experimental work [5] has shown that once a careful temperature stabilization is provided, a surprisingly good frequency stability is observed in semiconductor lasers. Their conclusions seem to suggest that once the network has been set up, only a very slow control or even an occasional check could be adequate to keep the transmitters at their nominal frequency. In STARNET, frequency adjustments could be triggered or controlled by the network protocol, and information on the stability of the comb could be exchanged through the packet network, so as to avoid critical conditions.

Finally, the first node in the comb could be locked to an absolute reference, distributed throughout the star. This reference needs to be *one spectral line only*, which would establish the 'origin' of the frequency comb.

This attractive picture can be preserved even if the chosen transmission format requires direct laser modulation as in FSK. In this case the fixed receiver would have a separate LO laser which would lock on the previous node. In addition, some circuitry would keep the node transmitter laser locked on the fixed receiver LO. This locking can be detected around baseband, in the fixed receiver. This solution is more complicated but it permits to avoid external modulators.

4.5 Implementation of BNC2

The BNC2 transmitter is identical to the BNC1 one. except for the fact that there are three TDM data streams instead of two, 'C', 'P1' and 'P2'. The 'C' stream receiver is identical. The additional receiver extracts the 'P1' stream from the longer wavelength adjacent node
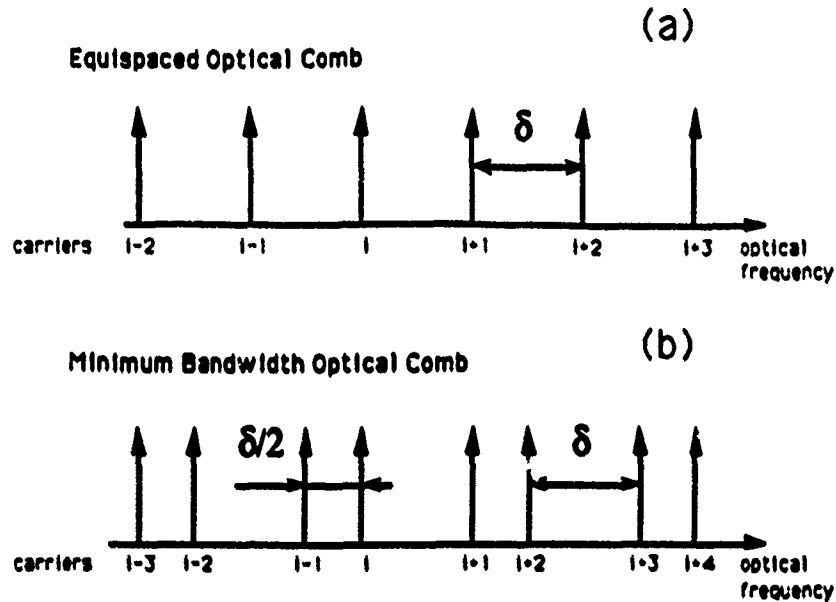
Figure 10: Equispaced and minimum optical bandwidth occupancy [2] allocation of carriers in the frequency comb, for heterodyne detection. The quantity $\delta/2$ is the minimum acceptable electrical channel spacing at the receiver IF.

and the 'P2' from the shorter wavelength one. Both are visible at IF in the fixed receiver (see Fig. 11). Frequency stabilization is again done with respect to the previous carrier in the frequency comb.

5 Data stream multiplexing

In the previous section we assumed that TDM was used to multiplex different data streams on the same carrier. In this section we explore some alternative solutions and compare them to TDM. The possible choices are:

- TDM;

- combined modulation formats;

- multilevel modulation.

We finally briefly discuss the maximum number of nodes that the network can support, on the basis of sensitivity considerations alone.

5.1 Performance of TDM

When TDM is used, the data rate of the transmitted bit stream is the sum of the bit rates of the 'P' and 'C' streams. Both the transmitter and the receiver in the node must be built to handle information at this speed.

For any given modulation format, the S/N ratio is inversely proportional to the transmission rate. The

penalty, in dB, on the 'C' link due to the multiplexing of the 'P' link is therefore:

$$\Delta = 10 \log_{10}\left(\frac{c+p}{c}\right) \qquad (1)$$

where c and p stand for the rates of the 'C' and 'P' links respectively. This way of dividing transmission power is very clean and no excess penalty occurs. A plot of Δ is shown in Fig. 12.

5.2 Combined modulation

Some modulation formats can be utilized simultaneously on the same carrier, and detected independently. For instance ASK is multiplexable with any modulation format which does not rely on transmission power variations (PSK, DPSK, FSK, POLSK, etc.). Obviously the modulation depth of the ASK stream cannot be large, since the other data stream needs power to support itself.

We computed the performance of the combined modulation approach for ASK with heterodyne PSK and DPSK. In both cases the ASK stream is extracted squaring the IF signal to get rid of the phase modulation. IF filtering has to be very loose to ensure complete phase modulation cancellation in the squarer. Instead post-detection filtering is as tight as possible. ASK modulation depth is adjusted so that, for a given carrier peak power, the error probability of ASK and the companion stream be equal. The results are plotted in Fig. 12. The
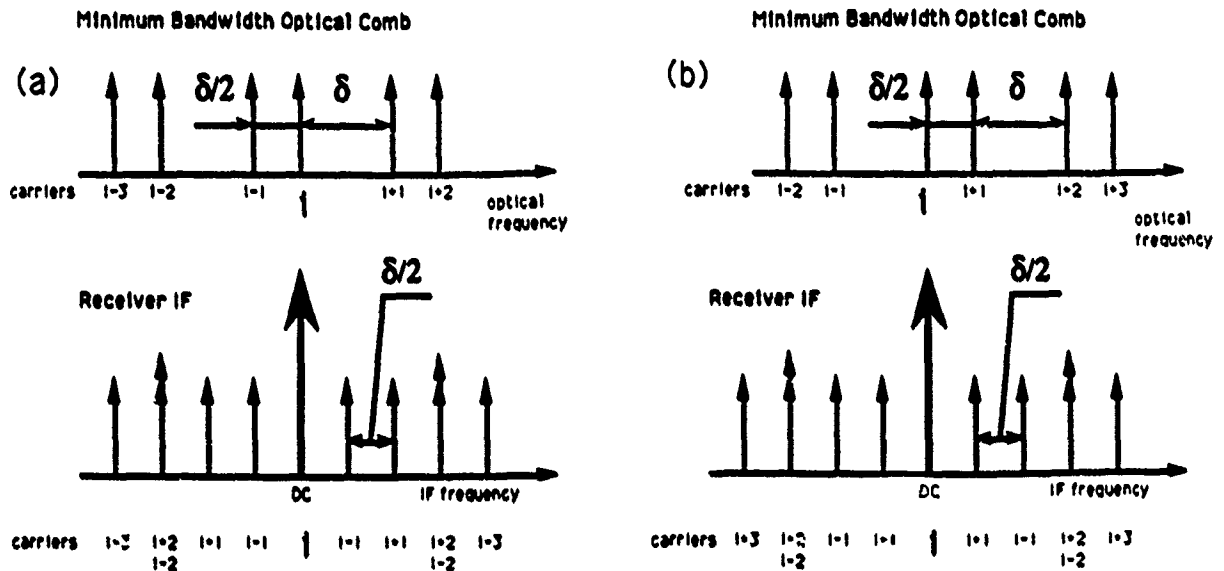
Figure 11: Carrier visibility at IF in the fixed receiver of node i, employing the node carrier as LO and minimum-optical-bandwidth allocation of carriers. In (a) and (b) the two possible IF spectra are reported, depending on the i-th node position in the comb.

abscissa is the speed ratio:

$$\rho = (\text{speed ASK})/(\text{speed companion}) \qquad (2)$$

and the ordinate is the penalty with respect to transmission of the companion stream alone.

The curves show a high penalty as long as $\rho < 5$. A logical choice is therefore that of using ASK to encode the 'P' stream and devote the other format to the higher speed 'C' link.

Also, this solution is more suitable for BNC1, since only one 'P' stream is to be supported. With BNC2 two 'P' streams are present, and an additional multiplexing strategy must be used over the ASK modulation to accomodate both.

Yet, combined modulation can be rather attractive, depending on the network parameters. We also emphasize that *single* LiNbO$_3$ modulators combining phase and amplitude modulation are feasible. Therefore no additional loss due to the use of two cascaded modulators occurs.

5.3 Multilevel transmission

In a multilevel transmission system, each transmitted symbol carries more than one information bit. If it is n bits, $n - k$ of these bits can be devoted to the 'C' link, and k of them to the 'P' link. This way, demultiplexing the different streams does not require a TDM frame extraction or an additional mo-demodulation.

In general, the use of multilevel transmission can be beneficial in a WDM network environment where bandwidth may be the limiting factor for the maximum number of nodes.

On the other hand, only two modulation formats seem to be mature for coherent optical multilevel transmission: PSK (and maybe QAM) and POLSK [17]. All these schemes are somewhat complicated and need experimental confirmation.

5.4 Maximum number of users

Assuming a DPSK implementation, the ideal sensitivity is 21 photons/bit. For a 3 Gbit/s link, this amounts to -50.9 dBm. For a thousand users, the equivalent link sensitivity is still -20.9 dBm. Since there are *two* receivers in each node, the incoming light must be split and therefore a 3dB degradation of sensitivity occurs. In addition, if the transmitter laser is used as LO for the fixed receiver, both the transmission power and the LO power are lower. Therefore the true system margin is certainly lower than the above figure suggests. Nevertheless, present day technology seems to be adequate to cope with these power levels. Depending on network design parameters, reasonable numbers for the maximum number of nodes could range between a few hundreds and several thousands.
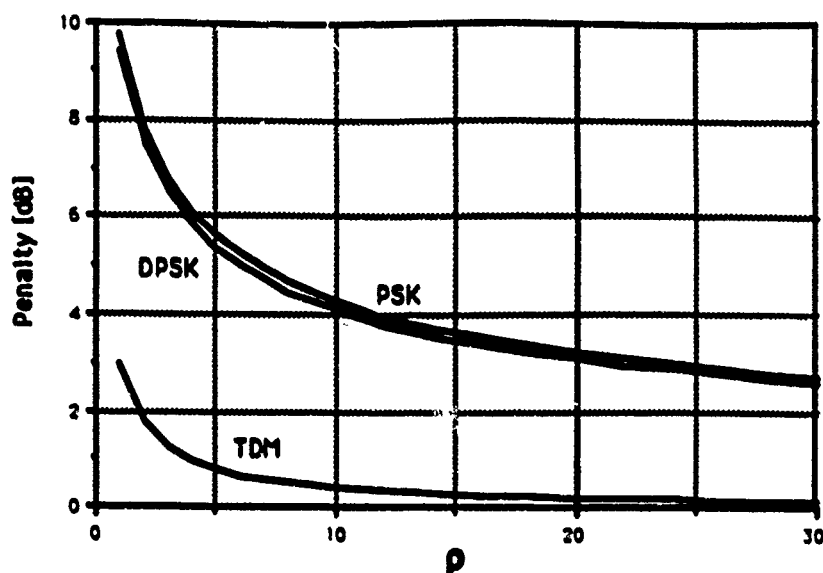
Figure 12: Penalty with respect to transmission of the 'C' stream alone, as a function of ρ = (speed 'P' stream)/(speed 'C' stream), computed for an error probability of 10^{-9} on both streams, for ASK-PSK, ASK-DPSK, TDM.

6 Enhanced topologies

In Section 3 we introduced the basic topologies that can be implemented for the packet network. Here we will explore some alternatives based on *the same elementary node structures* outlined in Section 3.

For a single ring arrangement, assuming one thousand nodes in the network, evenly distributed over an area of six kilometers in diameter, and an FDDI-like protocol with a per-node processing delay equal to the FDDI one, the worst-case packet latency time is 20.6 ms, almost entirely due to the propagation time. This figure does not account for queueing delays at the transmitting node. It can be be considered acceptable for many applications, and sufficient for conveying telephony on a virtual circuit basis. However, it might be inadequate for some time-sensitive services.

As far as capacity is concerned, there might be heavy flows of traffic over the ring which would involve just a few nodes. For instance a master workstations can provide system and storage support for a cluster of other workstations, or local mainframes could serve a number of graphic terminals. In the single ring arrangement this traffic would propagate throughout the whole network, despite the fact that it is logically confined to a few nodes, causing the network transmission capacity to be saturated.

In STARNET a solution is possible, which permits to improve both capacity and latency and does not require any changes in the optics or low-level electronics of the basic node configurations. The idea is similar to that of *bridging*, but in this context it helps to reduce packet delays too.

The solution for BNC1 is depicted in Fig. 13. We have already pointed out that to close the packet network ring, there must be a node which uses its high speed link to get the data from the last node in the ring and relay them along the chain. This station, which we now call 'Master Station' (M in Fig. 13), has its transmitter set at the first position in the comb of carriers of the ring.

We divide the original ring into several independent rings (henceforth 'subrings'). These subrings are contiguous in frequency and are logically connected as follows: the Master Station of the $(i+1)$-th subring, that we call M_{i+1}, detects the 'P' data stream circulating in the i-th subring (R_i) through its fixed receiver. Whenever a packet is detected, whose destination does not belong to R_i, it is picked up by M_{i+1} and relayed into R_{i+1}. M_{i+1} cannot extract this packet from the stream of R_i, but the packet is read by M_i too, which will check its 'subring destination' and take it out.

Many subrings can be cascaded this way. To complete the network a final 'back-bouncing' node is needed. We call it Master of Master stations (MM in Fig. 13), that takes the inter-subring traffic out of the last subring and relays it into the first subring.

The overall resulting topology is shown in Fig. 14. Each subring internal traffic does not propagate to the rest of the network. However, inter-subring traffic flows normally, as if no discontinuity existed.

Such a solution confers to the proposed network a high degree of flexibility, which is impossible with DQDB or FDDI[4]. This arrangement is easily reconfigurable, since any transmitter can be re-tuned to another frequency to change its subring assignment. Subrings do not need to have the same size. In this scenario a deep optimization can be performed and the effective capacity of the network can be greatly increased.

As for delays, the traffic confined within a subring experiences a propagation delay which is established by the number of nodes in *that* subring only. In our network proposal, therefore, bridging also cuts down delays. This permits to hook up to the same subring those nodes that need to talk faster among them.

However, there might be time-sensitive information which needs to travel across rings. If a packet is an inter-subring one, the worst-case maximum delay that it experiences is identical to that occuring in the unsegmented picture.

A better solution would be as follows: when M_j receives a packet from subring R_{j-1} whose destination is beyond subring R_j, it forwards it *directly* to R_{j+1}, instead of squeezing it into the data flow of R_j. If this were possible, the equivalent topology would be that of Fig. 15. In practice, an inner ring among the Master Stations would be formed, so that inter-subring traffic could rapidly reach the destination subring in a few hops.

This topology is realizable with either BNC1 or BNC2. It can be shown that with BNC1 it is necessary to use two nodes to make up a Master Station with the additional forwarding capability outlined above.

A more attractive implementation is possible with BNC2. With reference to Fig. 6, we assume that the leftmost node is the Master Station. As pointed out in Section 3, this node does not need to use its tunable receiver to close the subring, as opposed to the BNC1 subring implementation. Instead, the tunable receiver can be used to give rise to an inner packet network (Main Ring) that connects only the Master Stations, as shown in Figure 16. As a result the Master Stations use their 'P' streams to handle subring traffic and devote the capacity of their 'C' links to the Main Ring. This way, the data flows of the subrings and that of the Main Ring are completely isolated from each other and there is no need for the Master Stations to squeeze inter-subring traffic into the subring traffic flow. Another remarkable point is that the Main Ring could be implemented at a speed

different from that of the subring packet networks, being based on the use of part or all the capacity of the 'C' link. In particular it could be faster, enabling heavy inter-ring traffic to be exchanged without saturating the Main Ring. Only the Master Stations would be equipped with the faster electronics necessary to do so, increasing cost-effectiveness.

In the context of this enhanced topology, BNC2 turns out to be an appealing option, despite the more complicated node electronics. This overhead is paid for in terms of better network performance.

As for the number of Master Station nodes necessary to implement the above architectures, we remark that several tens of nodes, if not hundreds, could find place in each standard subring. The need of one Master Station per subring does not seem to be a heavy toll to pay.

7 High speed packet traffic

The solutions described so far provide effective answers to traffic which is either packetizable at low-to-medium speed, or conveyable at very high speed through circuit connections whose set-up time is relatively slow.

It can be shown that, using the packet network to exchange all the control information, the circuit interconnection set-up time is equal to the worst-case propagation and processing delay over the packet network, plus twice the packet queuing delay, plus the maximum tuning delay of one receiver laser source.

Recently, very high speed packetized interfaces have been proposed (for instance the HIPPI[5] interface), in the range of several hundred megabytes per second. The above delays would be completely inadequate if a single packet transmission involved the set-up and tearing-down of a high-speed link.

In the following we show that STARNET has the potential to accomodate this type of traffic too.

7.1 Packet ring subnetworks

Arbitrary subsets of nodes in STARNET can form logical rings among them through the 'C' links and the tunable receivers. This way it would be possible to realize ring subnetworks involving a limited number of nodes and working at up to the speed of the 'C' link. The main advantage is that these subnetworks do not interfere with each other and are independent from the underlying common packet network. Drawbacks are that per-node bandwidth is inversely proportional to the number of stations

---

[4] To achieve an effect similar to bridging, DQDB provides *erasure nodes* but the way they operate is less effective than the above scenario for STARNET.

[5] As of now, this interface is being standardized as a point-to-point high speed link, although the data format and transmission procedure is packet-compatible. In fact the target is that of extending the standard to a generalized packet network.

involved. Also, the connection or disconnection of a node causes interruption of the data flow for a time comparable with that needed for a node to tune to another arbitrary node.

Nevertheless, this configuration is attractive for clusters of powerful computers that must be connected in a stable way. Delays are very low since the number of nodes to be traversed is low. Packet protocols or fixed TDMA capacity allocation could be used.

7.2 Multihop subnetworks

Many recent proposals advocate the use of multihop logical topologies for optical networks (for instance ShuffleNet [6, 7]). These networks are thought to cover effectively the very high speed packet traffic segment, interconnecting a somewhat large number of high-speed users.

A binary shuffle net, or a binary multihop network, can be implemented among STARNET stations, provided that they either have an extra tunable receiver or they own two basic nodes. The second option is obvious, but it requires the use of two complete elementary nodes.

The first alternative would violate the BNC1 or BNC2 configurations. However, the extra receiver would not interfere with the main network modes of operation. It could be seen as an upgrade, that a potential user could add at a later time, when the related service becomes necessary. Having two tunable receivers at its disposal, every multihop network node could tune to two different nodes. At the same time, it could use its 'C' link to TDM multiplex two streams of data on its transmitter laser. These streams would in turn be read by two different nodes of the multihop network.

The presence of a second tunable receiver involves an additional penalty due to the fact that the incoming light must be divided among three receivers instead of two. Yet, the extra penalty amounts to only 1.8 dB.

This STARNET hybrid solution is attractive since not all the nodes in the network need to have the above additional capabilities, but only the ones that really need to exchange this particular type of traffic.

8 Broadcasting and video-conferencing

The backbone provided by the packet network to handle circuit interconnect control information can be used for a number of applications.

To set up a broadcast high speed transmission, control information about it (a broadcast 'call') can be sent over the packet network. Any node can tune independently on the broadcasting one.

Video-conferencing is one of the most attractive facilities for a future network to provide. A more appropriate term for what people have in mind is *remote meeting*: anybody must be able to join the meeting through his own workstation, on whose screen *all* the other participants appear.

Remote meeting can be handled through the packet network and the use of broadcasting for a low-speed compressed video format ($\approx$ 2Mb/s per video signal).

In this picture one node is *the broadcast station*. At the meeting set up, all the other participants are requested to tune to the broadcast station high speed link. A series of virtual circuits is also arranged on the packet network, linking each participant's terminal to the broadcast station. This station picks up all the data streams and TDM-multiplexes them onto its 'C' data stream, which is broadcast. Each node receives this flow of information that contains all the video signals in a multiplexed fashion, and displays them to the meeting attendee.

The above solution may turn out not to be adequate if either the packet network is not capable of supporting enough users, or because non-compressed or HDTV video quality is wanted. Another arrangement is possible, which completely relies on the high-speed links (except for control information).

The nodes involved in the meeting form a ring subnetwork using their high speed links. Each station is assigned a TDM slot and retransmits what it receives. It takes out its own old data and put new ones in its time slot. Synchronization is established by the flow of the different TDM frames.

As a result, at each station the complete visibility of all the video signals is achieved and no overhead is borne by the packet network. The connection limit is set by the capacity of the high speed link.

Broadcast quality HDTV has recently proved fesible at 44 Mbit/s. With reasonable transmission speeds as many as 10-15 attendees can be multiplexed. For standard (non-HDTV) video the number grows to 60-70.

9 STARNET experiment

An experimental STARNET prototype is currently being built at the Optical Communication Research Laboratory, Stanford University. It consists of four BNC1 nodes. The 'C' links are implemented through heterodyne PSK at 3Gbit/s. The 'P' stream is multiplexed by superimposing ASK modulation (see Section 5.2). The packet network is organized on a single ring and makes use of standard FDDI boards.

10 Conclusions

We propose a new broadband local area network, STAR-NET, based on a physical passive star topology. STAR-NET offers all users a moderate-speed packet network and a high-speed WDM circuit interconnect. STARNET can also provide an arbitrary subset of its nodes with a broadband packet multihop network. As a result, STAR-NET can support very diverse types of traffic in an optimal way.

We describe in detail the network structure and the options for implementation of the elementary nodes. The actual realisation of STARNET seems to be feasible with presently available technologies.

An effort toward an experimental demonstration of a 4-node, 3 Gbit/s per node, FDDI-compatible (at the packet network level) STARNET is currently in progress at the Optical Communication Research Laboratory, Stanford University.

References

[1] Chlamtac, I., Franta, W.: "Rationale, directions, and issues surrounding high speed networks"; Proceedings of the IEEE, vol. 78, no. 1, January 1990.

[2] O'Byrne, V.: "A method for reducing the channel spacing in a coherent optical heterodyne system"; IEEE Photonics Technology Letters, vol. 2, no. 7, July 1990.

[3] Mon-Song Chen, et al.: "A media access protocol for packet-switched wavelength division multiaccess metropolitan area networks"; IEEE J-SAC, vol. 8, no. 6, pp. 1048-1057, August 1990.

[4] Terumi Chikama, et al.: "Optical heterodyne image-rejection receiver for high-density optical frequency division multiplexing system"; IEEE JSAC, vol. 8, no. 6, pp. 1087-1094, August 1990.

[5] Sessa, w., B., et al.: "Multichannel frequency control for multilocation network applications"; OFC '91, Proceedings, paper WH5, p. 98, San Diego, February 19-22, 1991.

[6] Acampora, A., S.: "A multichannel multihop local lightwave network"; Proc. Globecom 1987, pp. 1500-1510.

[7] Hluchyj, M., G., Karol, M., J.: "ShuffleNet: an application of generalized perfect shuffles to multihop lightwave networks"; Proc. Infocom 1988.

[8] Habbab, I., M., I., et al.: "Protocols for very high speed optical fiber local area networks using a passive star topology"; IEEE JLT, vol. 5, no. 12, pp. 1782-1793, December 1987.

[9] Mehravari, N.: "Improved multiple access schemes for very high speed optical fiber local area networks using a passive star topology"; GLOBECOM '89, paper 8.2.1.

[10] Chlamtac, I., Gans, A.: "Design alternatives of asynchronous WDM star networks"; ICC '89.

[11] Gans, A., Koren, : "WDM passive stars - Protocols and performance analysis"; INFOCOM '91.

[12] Leung, , Takawira, : "FDM packet multiple access network using coherent fiber optic communication"; ICC '90.

[13] Ofek, Y., Sidi, M.: "Design and analysis of a hybrid access control to an optical star using WDM"; INFOCOM '91, paper 2A.3.1.

[14] Labourdette, J., F., Acampora, A., S.: "Wavelength agility in multihop lightwave networks"; INFOCOM '90, pp. 1022-1029.

[15] Labourdette, J., F., Acampora, A., S.: "Partially reconfigurable multihop lightwave networks"; GLOBECOM '90, paper 300.6.1.

[16] Bannister, J., A., et al.: "Topological design of wavelength-division optical networks"; INFOCOM '90, pp. 1005-1013.

[17] Benedetto, S., Poggiolini, P.: "Performance evaluation of multilevel polarisation shift keying modulation"; IEE Electronics Letters, vol. 26, no. 4, pp. 244-246, February 25th, 1990.
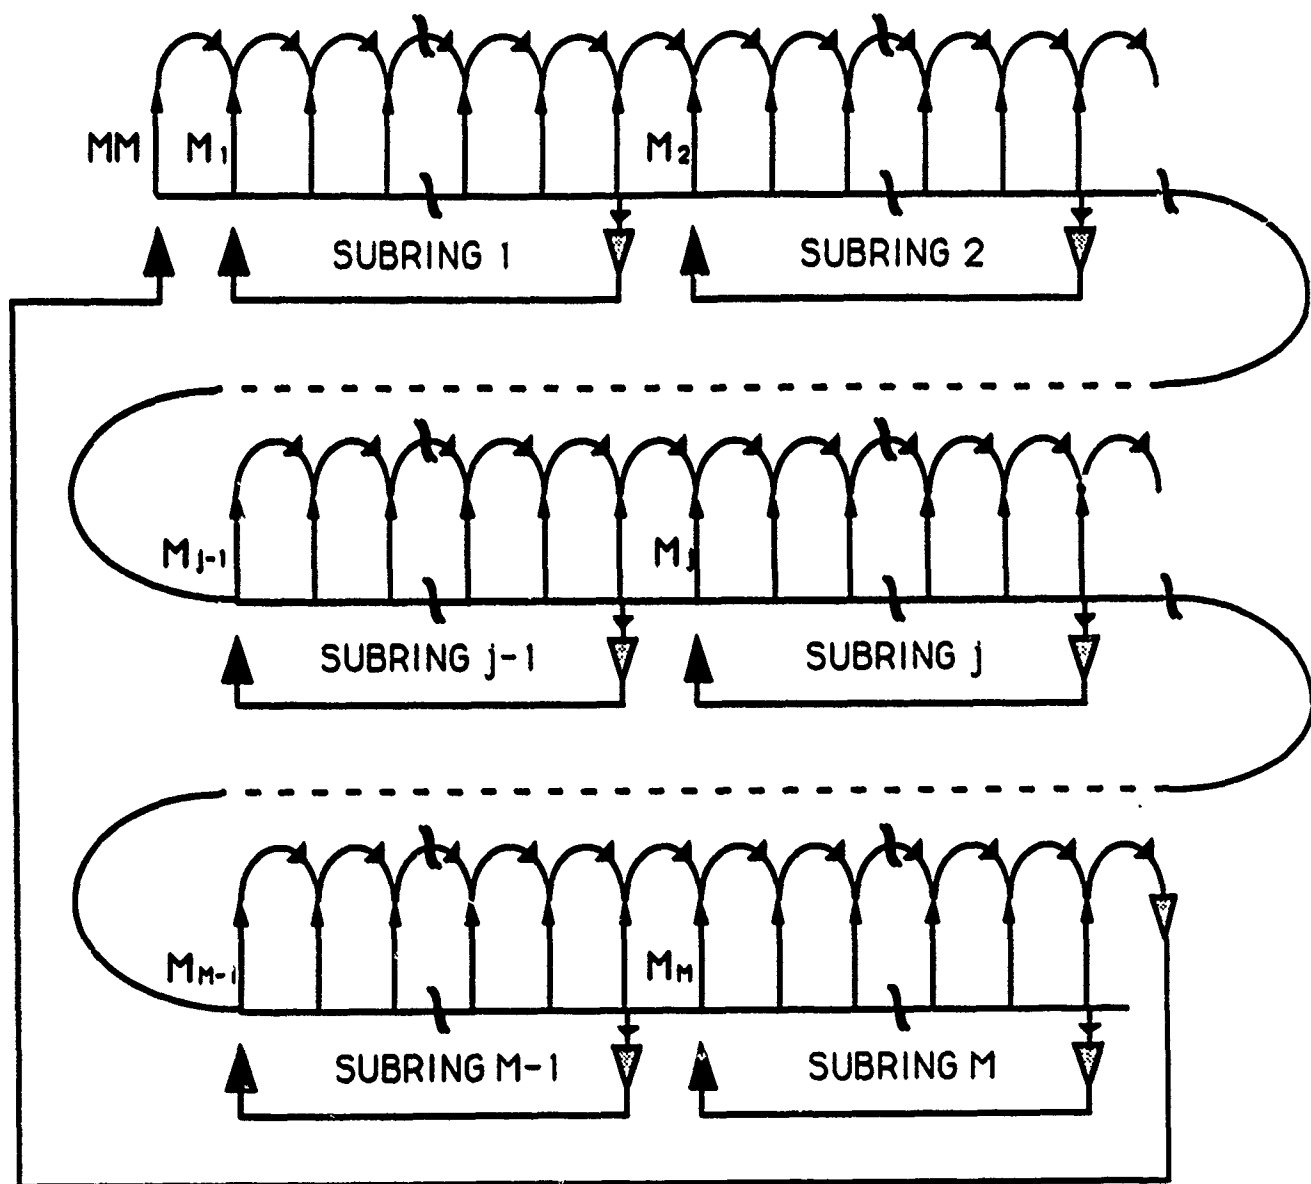
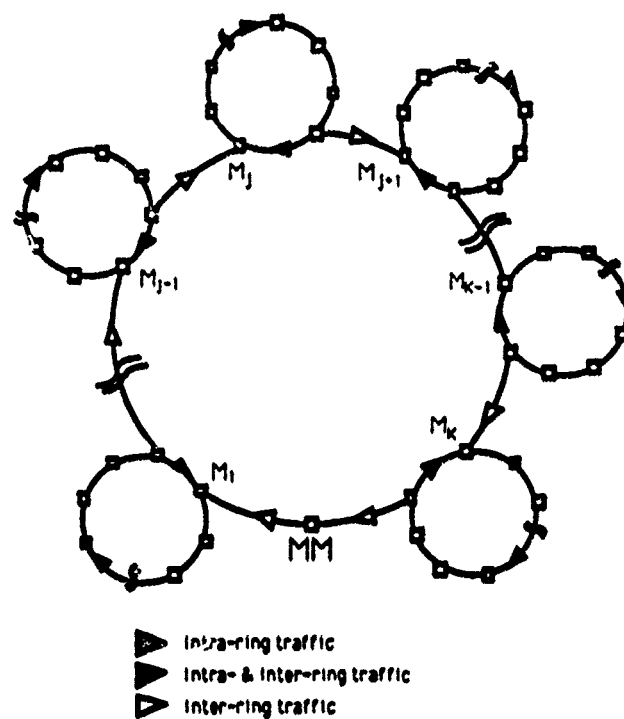Figure 13: Logical connection diagram in a segmented packet network, implemented with BNC1.

Intra-ring traffic
Intra- & inter-ring traffic
Inter-ring traffic

Figure 14: Data stream flow and virtual topology of the segmented packet network.

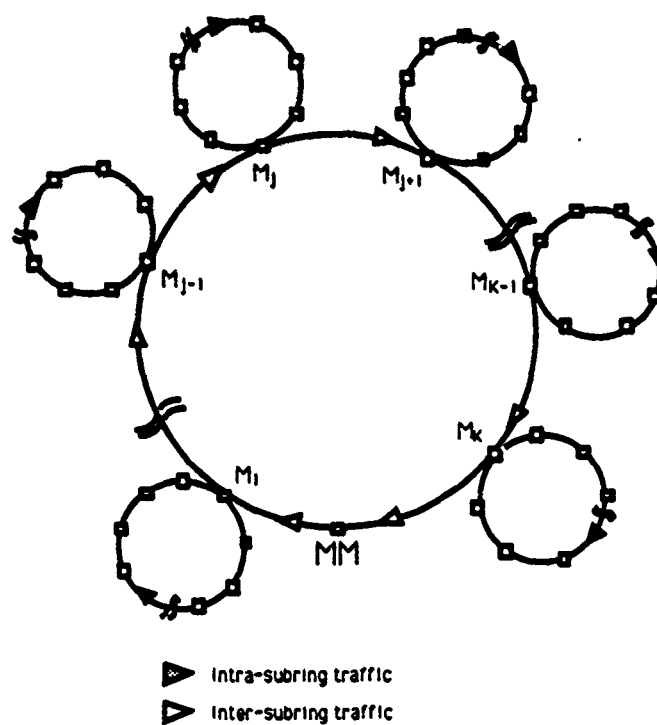Intra-subring traffic
Inter-subring traffic

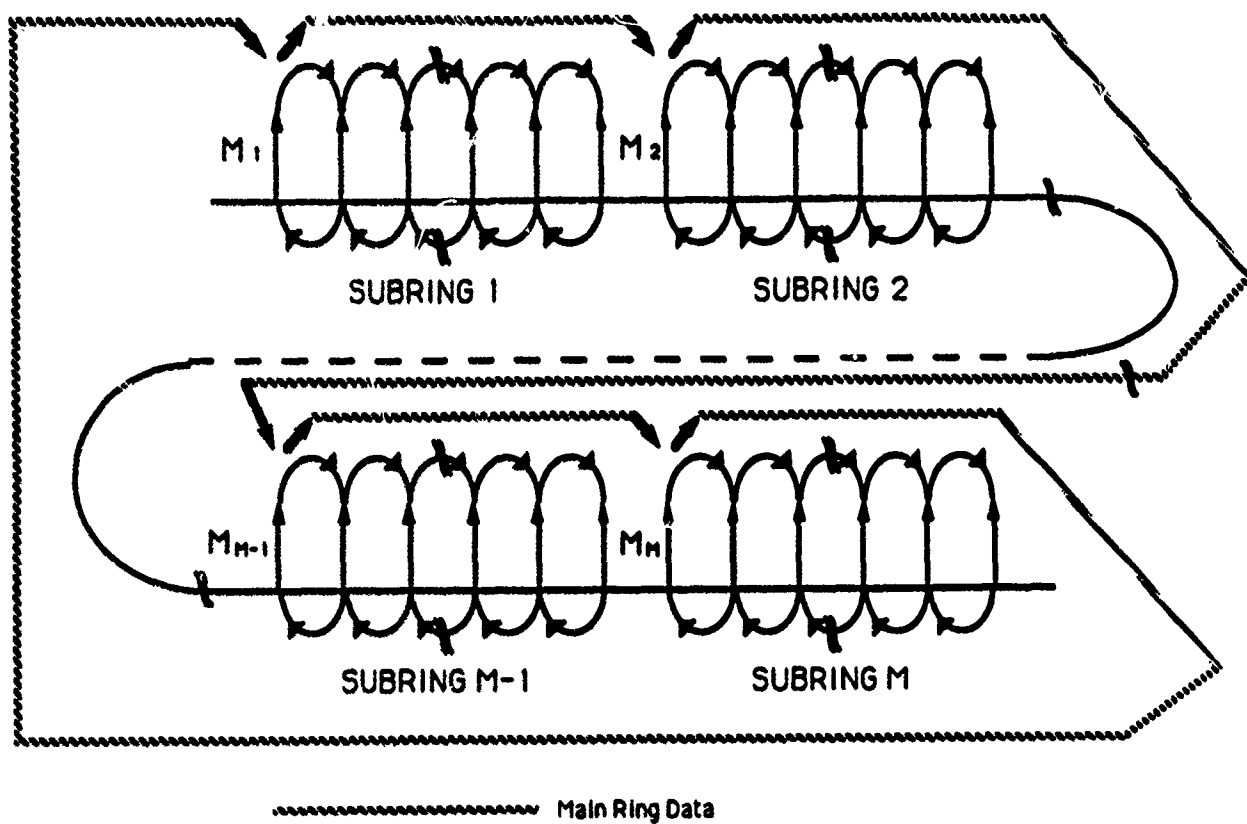Figure 15: Data stream flow and virtual topology of the ring-of-subrings packet network.

Figure 16: Logical connection diagram of the ring-of-subrings packet network, using BNC2.